

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

A System and Related Methods
for
Reducing Source Filter Invocation in a Development
Project

Inventor(s):
Daniel J. Miller
Eric H. Rudolph

ATTORNEY'S DOCKET NO. MS1-631US

1
2 **TECHNICAL FIELD**

3 This invention generally relates to processing media content and, more
4 particularly, to an interface and related methods for reducing source filter
5 invocation during execution of a development project.
6

7 **BACKGROUND**

8 Recent advances in computing power and related technology have fostered
9 the development of a new generation of powerful software applications. Gaming
10 applications, communications applications, and multimedia applications have
11 particularly benefited from increased processing power and clocking speeds.
12 Indeed, once the province of dedicated, specialty workstations, many personal
13 computing systems now have the capacity to receive, process and render
14 multimedia objects (e.g., audio and video content). While the ability to display
15 (receive, process and render) multimedia content has been around for a while, the
16 ability for a standard computing system to support true multimedia editing
17 applications is relatively new.

18 In an effort to satisfy this need, Microsoft Corporation introduced an
19 innovative development system supporting advanced user-defined multimedia
20 editing functions. An example of this architecture is presented in US Patent No.
21 5,913, 038 issued to Griffiths and commonly owned by the assignee of the present
22 invention, the disclosure of which is expressly incorporated herein by reference.

23 In the '038 patent, Griffiths introduced the an application program interface
24 which, when exposed to higher-level development applications, enable a user to
25 graphically construct a multimedia processing project by piecing together a

1 collection of "filters" exposed by the interface. The interface described therein is
2 referred to as a filter graph manager. The filter graph manager controls the data
3 structure of the filter graph and the way data moves through the filter graph. The
4 filter graph manager provides a set of software objects for communication
5 between a filter graph and its application. According to one implementation,
6 individual filters of a filter graph architecture are implemented as component
7 object model (COM) objects of an operating system executing on a computer,
8 each implementing one or more interfaces which contain a predefined set of
9 functions, called methods. Methods are called by an application program or other
10 component objects in order to communicate with the object exposing the interface.
11 The application program can also call methods or interfaces exposed by the filter
12 graph manager object.

13 Filter graphs work with data representing a variety of media (or non-media)
14 data types, each type characterized by a data stream that is processed by the filter
15 components comprising the filter graph. A filter positioned closer to the source of
16 the data is referred to as an upstream filter, while those further down the
17 processing chain is referred to as a downstream filter. For each data stream that
18 the filter handles it exposes at least one virtual pin (i.e., distinguished from a
19 physical pin such as one might find on an integrated circuit). A virtual pin can be
20 implemented as a COM object that represents a point of connection for a
21 unidirectional data stream on a filter. Input pins represent inputs and accept data
22 into the filter, while output pins represent outputs and provide data to other filters.
23 Each of the filters include at least one memory buffer, wherein communication of
24 the media stream between filters is accomplished by a series of "copy" operations
25 from one filter to another.

As introduced in Griffiths, a filter graph has three different types of filters: source filters, transform filters, and rendering filters. A source filter is used to load data from some source; a transform filter processes and passes data; and a rendering filter renders data to a hardware device or other locations (e.g., saved to a file, etc.). An example of a filter graph for a simplistic media rendering process is presented with reference to Fig. 1.

Fig. 1 graphically illustrates an example filter graph for rendering media content. As shown, the filter graph 100 is comprised of a plurality of filters 102-114, which read, process (e.g., transform) and render media content from a selected source file. As shown, the filter graph includes each of the types of filters described above, interconnected in a linear fashion.

Products utilizing the filter graph have been well received in the market as it has opened the door to multimedia editing using otherwise standard computing systems. It is to be appreciated, however, that the construction and implementation of the filter graphs are computationally intensive and expensive in terms of memory usage. Even the most simple of filter graphs requires and abundance of memory to facilitate the copy operations required to move data between filters. Thus, complex filter graphs can become unwieldy, due in part to the linear nature of conventional development system architecture. Moreover, it is to be appreciated that the filter graphs themselves consume memory resources, thereby compounding the issue introduced above.

Thus, what is required is a filter graph architecture which reduces the computational and memory resources required to support even the most complex of multimedia projects. Just such a solution is disclosed below.

SUMMARY

This invention concerns a system and related interfaces supporting the processing of media content. In accordance with one aspect of the present embodiment, a method of representing a development project is presented comprising identifying a plurality of sources comprising the development project, determining whether any of the sources are required simultaneously and, if not, dynamically generating a filter graph representation of the development project utilizing a segment filter to couple a source to multiple processing threads. In this way, the source and source filter, typically coupled to each of the multiple processing threads, is shared among multiple processing threads, reducing the computational and memory requirements necessary to support execution of the development project.

BRIEF DESCRIPTION OF THE DRAWINGS

The same reference numbers are used throughout the figures to reference like components and features.

Fig. 1 is a graphical representation of a conventional filter graph representing a user-defined development project.

Fig. 2 is a block diagram of a computing system incorporating the teachings of the described embodiment.

Fig. 3 is a block diagram of an example software architecture incorporating the teachings of the described embodiment.

Fig. 4 is a graphical illustration of an example software-enabled matrix switch, according to an exemplary embodiment.

Fig. 5 is a graphical representation of a data structure comprising a programming grid to selectively couple one or more of a scalable plurality of input pins to a scalable plurality of output pins of the matrix switch filter, in accordance with one aspect of the described embodiment.

Fig. 6 is a graphical illustration denoting shared buffer memory between filters, according to one aspect of the described embodiment.

Fig. 7 is a flow chart of an example method for generating a filter graph, in accordance with one aspect of the described embodiment.

Fig. 8 is a flow chart of an example method for negotiating buffer requirements between at least two adjacent filters, according to one aspect of the described embodiment.

Fig. 9 graphically illustrates an overview of a process that takes a user-defined editing project and composites a data structure that can be used to program the matrix switch.

Fig. 10 graphically illustrates the project of Fig. 9 in greater detail.

Fig. 11 shows an exemplary matrix switch dynamically generated in support of the project developed in Figs. 9 and 10, according to one described embodiment.

Fig. 12 illustrates a graphic representation of an exemplary data structure that represents the project of Fig. 10, according to one described embodiment.

Figs. 13-18 graphically illustrate various states of a matrix switch programming grid at select points in processing the project of Figs. 9 and 10 through the matrix switch, in accordance with one described embodiment.

Fig. 19 is a flow chart of an example method for processing media content, in accordance with one described embodiment.

1 Fig. 20 illustrates an example project with a transition and an effect, in
2 accordance with one described embodiment.

3 Fig. 21 shows an exemplary data structure in the form of a hierarchical tree
4 that represents the project of Fig. 20.

5 Figs. 22 and 23 graphically illustrate an example matrix switch
6 programming grid associated with the project of Fig. 20 at select points in time,
7 according to one described embodiment.

8 Fig. 24 shows an example matrix switch dynamically generated and
9 configured as the grid of Figs. 22 and 23 was being processed, in accordance with
10 one described embodiment.

11 Fig. 25 shows an exemplary project in accordance with one described
12 embodiment.

13 Fig. 26 graphically illustrates an example audio editing project, according
14 to one described embodiment.

15 Fig. 27 depicts an example matrix switch programming grid associated with
16 the project of Fig. 26.

17 Fig. 28 shows an example matrix switch dynamically generated and
18 configured in accordance with the programming grid of Fig. 27 to perform the
19 project of Fig. 26, according to one described embodiment.

20 Fig. 29 illustrates an exemplary media processing project incorporating
21 another media processing project as a composite, according to yet another
22 described embodiment.

23 Fig. 30 graphically illustrates an example data structure in the form of a
24 hierarchical tree structure that represents the project of Fig. 29.
25

1 Figs 31-36 graphically illustrate various matrix switch programming grid
2 states at select points in generating and configuring the matrix switch to
3 implement the media processing of Fig. 29.

4 Fig. 38 illustrates an example matrix switch suitable for use in the media
5 processing project of Fig. 29, according to one described embodiment.

6 Fig. 38a graphically illustrates an example data structure in the form of a
7 hierarchical tree structure that represents a project that is useful in understanding
8 composites in accordance with the described embodiments.

9 Fig. 39 is a flow diagram that describes steps in a method in accordance
10 with one described embodiment.

11 Fig. 40 illustrates an example method of generating a filter graph, in
12 accordance with one aspect of the present invention.

13 Fig. 41 graphically illustrates an example reuse list, according to one aspect
14 of the present invention.

15 Fig. 42 illustrates an example method for source combining in support of
16 the method introduced in Fig. 40, according to one embodiment of the present
17 invention.

18 Fig. 43 graphically illustrates a timeline representation of source combining
19 introduced in Fig. 42.

20 Fig. 44 illustrates a block diagram of an example render engine utilizing a
21 segment filter in a filter graph, in accordance with one aspect of the present
22 invention.

23 Fig. 45 illustrates a flow chart of an example method of generating a filter
24 graph to reuse source filters, in accordance with one aspect of the present
25 invention.

Fig. 46 illustrates a flow chart of an example method of executing a development project utilizing a segment filter, in accordance with one aspect of the present invention.

DETAILED DESCRIPTION

Related Applications

This application is related to the following commonly-filed U.S. Patent Applications, all of which are commonly assigned to Microsoft Corp., the disclosures of which are incorporated by reference herein:

- Application Serial No. _____, entitled "An Interface and Related Methods for Reducing Source Accesses in a Development System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-643US;
- Application Serial No. _____, entitled "A System and Related Interfaces Supporting the Processing of Media Content", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-629US;
- Application Serial No. _____, entitled "A System and Related Methods for Reducing Memory Requirements of a Media Processing System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-632US;
- Application Serial No. _____, entitled "A System and Related Methods for Reducing the Instances of Source Files in a Filter Graph", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-633US;
- Application Serial No. _____, entitled "An Interface and Related Methods for Dynamically Generating a Filter Graph in a Development System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-634US;
- Application Serial No. _____, entitled "A System and Related Methods for Processing Audio Content in a Filter Graph", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-639US;
- Application Serial No. _____, entitled "A System and Methods for Generating an Managing Filter Strings in a Filter

Graph", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-642US;

- Application Serial No. _____, entitled "Methods and Systems for Processing Media Content", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-640US;
- Application Serial No. _____, entitled "Systems for Managing Multiple Inputs and Methods and Systems for Processing Media Content ", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-635US;
- Application Serial No. _____, entitled "Methods and Systems for Implementing Dynamic Properties on Objects that Support Only Static Properties", naming Daniel J. Miller and David Maymudes as inventors, and bearing attorney docket number MS1-638US;
- Application Serial No. _____, entitled "Methods and Systems for Efficiently Processing Compressed and Uncompressed Media Content", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-630US;
- Application Serial No. _____, entitled "Methods and Systems for Effecting Video Transitions Represented By Bitmaps", naming Daniel J. Miller and David Maymudes as inventors, and bearing attorney docket number MS1-637US;
- Application Serial No. _____, entitled "Methods and Systems for Mixing Digital Audio Signals", naming Eric H. Rudolph as inventor, and bearing attorney docket number MS1-636US; and
- Application Serial No. _____, entitled "Methods and Systems for Processing Multi-media Editing Projects", naming Eric H. Rudolph as inventor, and bearing attorney docket number MS1-641US.

Various described embodiments concern an application program interface associated with a development system. According to one example implementation, the interface is exposed to a media processing application to enable a user to dynamically generate complex media processing tasks, e.g., editing projects. In the discussion herein, aspects of the invention are developed within the general context of computer-executable instructions, such as program modules, being executed by one or more conventional computers. Generally,

1 program modules include routines, programs, objects, components, data structures,
2 etc. that perform particular tasks or implement particular abstract data types.
3 Moreover, those skilled in the art will appreciate that the invention may be
4 practiced with other computer system configurations, including hand-held devices,
5 personal digital assistants, multiprocessor systems, microprocessor-based or
6 programmable consumer electronics, network PCs, minicomputers, mainframe
7 computers, and the like. In a distributed computer environment, program modules
8 may be located in both local and remote memory storage devices. It is noted,
9 however, that modification to the architecture and methods described herein may
10 well be made without deviating from spirit and scope of the present invention.
11 Moreover, although developed within the context of a media processing system
12 paradigm, those skilled in the art will appreciate, from the discussion to follow,
13 that the application program interface may well be applied to other development
14 system implementations. Thus, the media processing system described below is
15 but one illustrative implementation of a broader inventive concept.

16 17 **Example System Architecture**

18 **Fig. 2** illustrates an example of a suitable computing environment 200 on
19 which the system and related methods for processing media content may be
20 implemented.

21 It is to be appreciated that computing environment 200 is only one example
22 of a suitable computing environment and is not intended to suggest any limitation
23 as to the scope of use or functionality of the media processing system. Neither
24 should the computing environment 200 be interpreted as having any dependency
25

1 or requirement relating to any one or combination of components illustrated in the
2 exemplary computing environment 200.

3 The media processing system is operational with numerous other general
4 purpose or special purpose computing system environments or configurations.
5 Examples of well known computing systems, environments, and/or configurations
6 that may be suitable for use with the media processing system include, but are not
7 limited to, personal computers, server computers, thin clients, thick clients, hand-
8 held or laptop devices, multiprocessor systems, microprocessor-based systems, set
9 top boxes, programmable consumer electronics, network PCs, minicomputers,
10 mainframe computers, distributed computing environments that include any of the
11 above systems or devices, and the like.

12 In certain implementations, the system and related methods for processing
13 media content may well be described in the general context of computer-
14 executable instructions, such as program modules, being executed by a computer.
15 Generally, program modules include routines, programs, objects, components,
16 data structures, etc. that perform particular tasks or implement particular abstract
17 data types. The media processing system may also be practiced in distributed
18 computing environments where tasks are performed by remote processing devices
19 that are linked through a communications network. In a distributed computing
20 environment, program modules may be located in both local and remote computer
21 storage media including memory storage devices.

22 In accordance with the illustrated example embodiment of Fig. 2 computing
23 system 200 is shown comprising one or more processors or processing units 202, a
24 system memory 204, and a bus 206 that couples various system components
25 including the system memory 204 to the processor 202.

Bus 206 is intended to represent one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnects (PCI) bus also known as Mezzanine bus.

Computer 200 typically includes a variety of computer readable media. Such media may be any available media that is locally and/or remotely accessible by computer 200, and it includes both volatile and non-volatile media, removable and non-removable media.

In Fig. 2, the system memory 204 includes computer readable media in the form of volatile, such as random access memory (RAM) 210, and/or non-volatile memory, such as read only memory (ROM) 208. A basic input/output system (BIOS) 212, containing the basic routines that help to transfer information between elements within computer 200, such as during start-up, is stored in ROM 208. RAM 210 typically contains data and/or program modules that are immediately accessible to and/or presently be operated on by processing unit(s) 202.

Computer 200 may further include other removable/non-removable, volatile/non-volatile computer storage media. By way of example only, Fig. 2 illustrates a hard disk drive 228 for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"), a magnetic disk drive 230 for reading from and writing to a removable, non-volatile

1 magnetic disk 232 (e.g., a "floppy disk"), and an optical disk drive 234 for reading
2 from or writing to a removable, non-volatile optical disk 236 such as a CD-ROM,
3 DVD-ROM or other optical media. The hard disk drive 228, magnetic disk drive
4 230, and optical disk drive 234 are each connected to bus 206 by one or more
5 interfaces 226.

6 The drives and their associated computer-readable media provide
7 nonvolatile storage of computer readable instructions, data structures, program
8 modules, and other data for computer 200. Although the exemplary environment
9 described herein employs a hard disk 228, a removable magnetic disk 232 and a
10 removable optical disk 236, it should be appreciated by those skilled in the art that
11 other types of computer readable media which can store data that is accessible by a
12 computer, such as magnetic cassettes, flash memory cards, digital video disks,
13 random access memories (RAMs), read only memories (ROM), and the like, may
14 also be used in the exemplary operating environment.

15 A number of program modules may be stored on the hard disk 228,
16 magnetic disk 232, optical disk 236, ROM 208, or RAM 210, including, by way of
17 example, and not limitation, an operating system 214, one or more application
18 programs 216 (e.g., multimedia application program 224), other program modules
19 218, and program data 220. In accordance with the illustrated example
20 embodiment of Fig. 2, operating system 214 includes an application program
21 interface embodied as a render engine 222. As will be developed more fully
22 below, render engine 222 is exposed to higher-level applications (e.g., 216) to
23 automatically assemble filter graphs in support of user-defined development
24 projects, e.g., media processing projects. Unlike conventional media processing
25 systems, however, render engine 222 utilizes a scalable, dynamically

reconfigurable matrix switch to reduce filter graph complexity, thereby reducing the computational and memory resources required to complete a development project. Various aspects of the innovative media processing system represented by a computer 200 implementing the innovative render engine 222 will be developed further, below.

Continuing with Fig. 2, a user may enter commands and information into computer 200 through input devices such as keyboard 238 and pointing device 240 (such as a "mouse"). Other input devices may include a audio/video input device(s) 253, a microphone, joystick, game pad, satellite dish, serial port, scanner, or the like (not shown). These and other input devices are connected to the processing unit(s) 202 through input interface(s) 242 that is coupled to bus 206, but may be connected by other interface and bus structures, such as a parallel port, game port, or a universal serial bus (USB).

A monitor 256 or other type of display device is also connected to bus 206 via an interface, such as a video adapter 244. In addition to the monitor, personal computers typically include other peripheral output devices (not shown), such as speakers and printers, which may be connected through output peripheral interface 246.

Computer 200 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 250. Remote computer 250 may include many or all of the elements and features described herein relative to computer 200 including, for example, render engine 222 and one or more development applications 216 utilizing the resources of render engine 222.

As shown in Fig. 2, computing system 200 is communicatively coupled to remote devices (e.g., remote computer 250) through a local area network (LAN) 251 and a general wide area network (WAN) 252. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet.

When used in a LAN networking environment, the computer 200 is connected to LAN 251 through a suitable network interface or adapter 248. When used in a WAN networking environment, the computer 200 typically includes a modem 254 or other means for establishing communications over the WAN 252. The modem 254, which may be internal or external, may be connected to the system bus 206 via the user input interface 242, or other appropriate mechanism.

In a networked environment, program modules depicted relative to the personal computer 200, or portions thereof, may be stored in a remote memory storage device. By way of example, and not limitation, Fig. 2 illustrates remote application programs 216 as residing on a memory device of remote computer 250. It will be appreciated that the network connections shown and described are exemplary and other means of establishing a communications link between the computers may be used.

Turning next to **Fig. 3**, a block diagram of an example development system architecture is presented, in accordance with one embodiment of the present invention. In accordance with the illustrated example embodiment of Fig. 3, development system 300 is shown comprising one or more application program(s) 216 coupled to render engine 222 via an appropriate communications interface 302. As used herein, application program(s) 216 are intended to represent any of a

1 wide variety of applications which may benefit from use of render engine 222
2 such as, for example a media processing application 224.

3 The communications interface 302 is intended to represent any of a number
4 of alternate interfaces used by operating systems to expose application program
5 interface(s) to applications. According to one example implementation, interface
6 302 is a component object model (COM) interface, as used by operating systems
7 offered by Microsoft Corporation. As introduced above, COM interface 302
8 provides a means by which the features of the render engine 222, to be described
9 more fully below, are exposed to an application program 216.

10 In accordance with the illustrated example implementation of Fig. 3, render
11 engine 222 is presented comprising source filter(s) 304A-N, transform filter(s)
12 306A-N and render filter 310, coupled together utilizing virtual pins to facilitate a
13 user-defined media processing project. According to one implementation, the
14 filters of system 300 are similar to the filters exposed in conventional media
15 processing systems. According to one implementation, however, filters are not
16 coupled via such interface pins. Rather, alternate implementations are envisioned
17 wherein individual filters (implemented as objects) make calls to other objects,
18 under the control of the render engine 222, for the desired input. Unlike
19 conventional systems, however, render engine 222 exposes a scalable, dynamically
20 reconfigurable matrix switch filter 308, automatically generated and dynamically
21 configured by render engine 222 to reduce the computational and memory
22 resource requirements often associated with development projects. As introduced
23 above, the pins (input and/or output) are application interface(s) designed to
24 communicatively couple other objects (e.g., filters).

1 In accordance with the example implementation of a media processing
2 system, an application communicates with an instance of render engine 222 when
3 the application 216 wants to process streaming media content. Render engine 222
4 selectively invokes and controls an instance of filter graph manager (not shown) to
5 automatically create a filter graph by invoking the appropriate filters (e.g., source,
6 transform and rendering). As introduced above, the communication of media
7 content between filters is achieved by either (1) coupling virtual output pins of one
8 filter to the virtual input pins of requesting filter; or (2) by scheduling object calls
9 between appropriate filters to communicate the requested information. As shown,
10 source filter 304 receives streaming data from the invoking application or an
11 external source (not shown). It is to be appreciated that the streaming data can be
12 obtained from a file on a disk, a network, a satellite feed, an Internet server, a
13 video cassette recorder, or other source of media content. As introduced above,
14 transform filter(s) 306 take the media content and processes it in some manner,
15 before passing it along to render filter 310. As used herein, transform filter(s) 306
16 are intended to represent a wide variety of processing methods or applications that
17 can be performed on media content. In this regard, transform filter(s) 306 may
18 well include a splitter, a decoder, a sizing filter, a transition filter, an effects filter,
19 and the like. The function of each of these filters is described more fully in the
20 Griffiths application, introduced above, and generally incorporated herein by
21 reference. The transition filter, as used herein, is utilized by render engine 222 to
22 transition the rendered output from a first source to a second source. The effect
23 filter is selectively invoked to introduce a particular effect (e.g., fade, wipe, audio
24 distortion, etc.) to a media stream.

1 In accordance with one aspect of the embodiment, to be described more
2 fully below, matrix switch filter 308 selectively passes media content from one or
3 more of a scalable plurality of input(s) to a scalable plurality of output(s).
4 Moreover, matrix switch 308 also supports implementation of a cascaded
5 architecture utilizing feedback paths, i.e., wherein transform filters 306B, 306C,
6 etc. coupled to the output of matrix switch 308 are dynamically coupled to one or
7 more of the scalable plurality of matrix switch input(s). An example of this
8 cascaded filter graph architecture is introduced in Fig. 3, and further explained in
9 example implementations, below.

10 Typically, media processed through source, transform and matrix switch
11 filters are ultimately passed to render filter 310, which provides the necessary
12 interface to a hardware device, or other location that accepts the renderer output
13 format, such as a memory or disk file, or a rendering device.

14 **Fig. 4** is a graphical illustration of an example software-enabled matrix
15 switch 308, according to one example embodiment of the present invention. As
16 shown, the matrix switch 308 is comprised of a scalable plurality of input(s) 402
17 and a scalable plurality of output(s) 404, wherein any one or more of the input(s)
18 402 may be iteratively coupled to any one or more of the output(s) 404, based on
19 the content of the matrix switch programming grid 406, automatically generated
20 by render engine 222. According to an alternate implementation introduced
21 above, switch matrix 308 is programmed by render engine 222 to dynamically
22 generate object calls to communicate media content between filters. In addition,
23 according to one implementation, matrix switch 308 includes a plurality of
24 input/output (I/O) buffers 408, as well as means for maintaining source, or media
25 time 410 and/or timeline, or project time 412. It is to be appreciated, however,

that in alternate implementations matrix switch 308 does not maintain both source and project times, relying on an upstream filter to convert between these times. As will be developed more fully below, matrix switch 308 dynamically couples one or more of the scalable plurality of inputs 402 to one or more of the scalable plurality of outputs 404 based, at least in part, on the media time 410 and/or the project time 412 and the content of matrix switch programming grid 406. In this regard, matrix switch 308 may be characterized as time-aware, supporting such advanced editing features as searching/seeking to a particular point (e.g., media time) in the media content, facilitating an innovative buffering process utilizing I/O buffers 408 to facilitate look-ahead processing of media content, and the like. Thus, it will be appreciated given the discussion to follow that introduction of the matrix switch 308 provides a user with an editing flexibility that was heretofore unavailable in a personal computer-based media processing system.

As introduced above, the inputs 402 and outputs 404 of matrix switch 308 are interfaces which facilitate the time-sensitive routing of data (e.g., media content) in accordance with a user-defined development project. Matrix switch 308 has a scalable plurality of inputs 402 and outputs 404, meaning that the number of inputs 402 and outputs 404 are individually generated to satisfy a given editing project. Insofar as each of the inputs/outputs (I/O) has an associated transfer buffer (preferably shared with an adjacent filter) to communicate media content, the scalability of the input/output serves to reduce the overall buffer memory consumed by an editing project. According to one implementation, output 1 is generally reserved as a primary output, e.g., coupled to a rendering filter (not shown).

According to one implementation, for each input 402 and output 404, matrix switch 308 attempts to be the allocator, or manager of the buffer associated with the I/O(s) shared with adjacent filters. One reason is to ensure that all of the buffers are of the same size and share common attributes so that a buffer associated with any input 402 may be shared with any output 404, thereby reducing the need to copy memory contents between individual buffers associated with such inputs/outputs. If matrix switch 308 cannot be an allocator for a given output (404), communication from an input (402) to that output is performed using a conventional memory copy operation between the individual buffers associated with the select input/output.

As introduced above, the matrix switch programming grid 406 is dynamically generated by render engine 222 based, at least in part, on the user-defined development project. As will be developed below, render engine 222 invokes an instance of filter graph manager to assemble a tree structure of an editing project, noting dependencies between source, filters and time to dynamically generate the programming grid 406. A data structure comprising an example programming grid 406 is introduced with reference to Fig. 5, below.

Turning briefly to **Fig. 5**, a graphical representation of a data structure comprising an example programming grid 406 is presented, in accordance with one embodiment of the present invention. In accordance with the illustrated example embodiment of Fig. 5, programming grid 406 is depicted as a two-dimensional data structure comprising a column along the y-axis 502 of the grid denoting input pins associated with a content chain (e.g., series of filters to process media content) of the development project. The top row along the x-axis 504 of the data structure denotes project time. With these grid "borders", the body 506 of

1 the grid 406 is populated with output pin assignments, denoting which input pin is
2 coupled to which output pin during execution of the development project. In this
3 way, render engine 222 dynamically generates and facilitates matrix switch 308.
4 Those skilled in the art will appreciate, however, that data structures of greater or
5 lesser complexity may well be used in support of the programming grid 406
6 without deviating from the spirit and scope of the present invention.

7 Returning to Fig. 4, matrix switch 308 is also depicted with a plurality of
8 input/output buffers 408, shared among all of the input(s)/output(s) (402, 404) to
9 facilitate advanced processing features. That is, while not required to implement
10 the core features of matrix switch 308, I/O buffers 408 facilitate a number of
11 innovative performance enhancing features to improve the performance (or at least
12 the user's perception of performance) of the processing system, thereby providing
13 an improved user experience. According to one implementation, I/O buffers 408
14 are separate from the buffers assigned to each individual input and output pin in
15 support of communication through the switch. According to one implementation,
16 I/O buffers 408 are primarily used to foster look-ahead processing of the project.
17 Assume, for example, that a large portion of the media processing project required
18 only 50% of the available processing power, while some smaller portion required
19 150% of the available processing power. Implementation of the shared I/O buffers
20 408 enable filter graph manager to execute tasks ahead of schedule and buffer this
21 content in the shared I/O buffers 408 until required. Thus, when execution of the
22 filter graph reaches a point where more than 100% of the available processing
23 power is required, the processing system can continue to supply content from the
24 I/O buffers 408, while the system completes execution of the CPU-intensive tasks.
25 If enough shared buffer space is provided, the user should never know that some

tasks were not performed in real-time. According to one implementation, shared buffers 408 are dynamically split into two groups by render engine 222, a first group supports the input(s) 402, while a second (often smaller) group is used in support of a primary output (e.g., output pin 1) to facilitate a second, independent output processing thread. The use of an independent output buffers the render engine from processing delays that might occur in upstream and/or downstream filters, as discussed above. It will be appreciated by those skilled in the art that such that matrix switch 308 and the foregoing described architecture beneficially suited to support media streaming applications.

As introduced above, the filter graph is time-aware in the sense that media (source) time and project execution time are maintained. According to one implementation, matrix switch 308 maintains at least the project clock, while an upstream filter maintains the source time, converting between source and project time for all downstream filters (i.e., including the matrix switch 308). According to one implementation, the frame rate converter filter of a filter graph is responsible for converting source time to project time, and vice versa, i.e., supporting random seeks, etc. Alternatively, matrix switch 308 utilizes an integrated set of clock(s) to independently maintain project and media times.

Having introduced the architectural and operational elements of matrix switch filter 308, **Fig. 6** graphically illustrates an example filter graph implementation incorporating the innovative matrix switch 308. In accordance with the illustrated example embodiment, filter graph 600 is generated by render engine 222 in response to a user defined development project. Unlike the lengthy linear filter graphs typical of convention development systems however, filter graph 600 is shown incorporating a matrix switch filter 308 to recursively route

the pre-processed content (e.g., through filters 602, 606, 610, 614 and 618, described more fully below) through a user-defined number of transform filters including, for example, transition filter(s) 620 and effects filter(s) 622. Moreover, as will be developed more fully below, the scalable nature of matrix switch filter 308 facilitates such iterative processing for any number of content threads, tracks or compositions.

According to one implementation, a matrix switch filter 308 can only process one type of media content, of the same size and at the same frame-rate (video) or modulation type/schema (audio). Thus, Fig. 6 is depicted comprising pre-processing filters with a parser filter 606 to separate, independent content type(s) (e.g., audio content and video content), wherein one of the media types would be processed along a different path including a separate instance of matrix switch 308. Thus, in accordance with the illustrated example embodiment of a media processing system, processing multimedia content including audio and video would utilize two (2) matrix switch filters 308, one dedicated to audio processing (not shown) and one dedicated to video processing. That is not to say, however, that multiple switch filters 308 could not be used (e.g., two each for audio and video) for each content type in alternate implementations. Similarly, it is anticipated that in alternate implementations a matrix switch 308 that accepts multiple media types could well be used without deviating from the spirit and scope of the present invention.

In addition filter graph 600 includes a decoder filter 610 to decode the media content. Resize filter 614 is employed when matrix switch 308 is to receive content from multiple sources, ensuring that the size of the received content is the same, regardless of the source. According to one implementation, resize filter 614

is selectively employed in video processing paths to adjust the media size of content from one or more sources to a user-defined level. Alternatively, resizer filter 614 adjusts the media size to the largest size provided by any one or more media sources. That is, if, for example, render engine 222 identifies the largest required media size (e.g., 1270x1040 video pixels per frame) and, for any content source not providing content at this size, the content is modified (e.g., stretched, packed, etc.) to fill this size requirement. The frame rate converter (FRC) and pack filter 618, introduced above, ensures that video content from the multiple sources is arriving at the same frame rate, e.g., ten (10) frames per second. As introduced above, the FRC also maintains the distinction between source time and project time.

In accordance with one aspect of the present invention, filter graph 600 is depicted utilizing a single, negotiated buffer 604, 608, 612, 616, etc. between adjacent filters. In this regard, render engine 222 reduces the buffer memory requirements in support of a development project.

From the point of pre-processing (filters 602, 606, 610, 614, 618), rather than continue a linear filter graph incorporating all of the transition 620 and effect 622 filter(s), render engine 222 utilizes a cascade architecture, recursively passing media content through the matrix switch 308 to apply to the transform filter(s) (e.g., 620, 622, etc.) to complete the execution of the development project. It will be appreciated by those skilled in the art that the ability to recursively pass media content through one or more effect and/or transition filters provided by the matrix switch filter 308 greatly reduces the perceived complexity of otherwise large filter graphs, while reducing memory and computational overhead.

Turning to **Fig. 7**, a flow chart of an example method for generating a filter graph is presented, in accordance with one aspect of the present invention. The method 700 begins with block 702 wherein render engine 222 receives an indication to generate a filter graph representing a user-defined development project (e.g., a media editing project). According to one example implementation, the indication is received from an application 224 via COM interface(s) 302.

In block 704, render engine 222 facilitates generation of the editing project, identifying the number and type of media sources selected by the user. In block 706, based at least in part on the number and/or type of media sources, filter graph manger 222 exposes source, transform and rendering filter(s) to effect a user defined media processing project, while beginning to establish a programming grid 406 for the matrix switch filter 308.

In block 708, reflecting user editing instructions, render engine 222 completes the programming grid 406 for matrix switch 308, identifying which inputs 402 are to be coupled to which outputs 404 at particular project times.

Based, at least in part, on the programming grid 406 render engine 222 generates a matrix switch filter 308 with an appropriate number of input 402 and output 404 pins to effect the project, and assembles the filter graph, block 710.

In block 712, to reduce the buffer memory requirements for the processing project, the render engine 222 instructs the filters populating the filter graph to (re)negotiate buffer memory requirements between filters. That is, adjacent filters attempt to negotiate a size and attribute standard so that a single buffer can be utilized to couple each an output pin of one filter to an input pin of a downstream filter. An example implementation of the buffer negotiation process of block 712 is presented in greater detail with reference to Fig. 8.

Turning briefly to Fig. 8, an example method of negotiating buffer requirements between adjacent filters is presented, in accordance with one example implementation of the present invention. Once the final connection is established to matrix switch 308, matrix switch 308 identifies the maximum buffer requirements for any filter coupled to any of its pins (input 402 and/or output 404), block 802. According to one implementation, the maximum buffer requirements are defined as the lowest common multiple of buffer alignment requirements, and the maximum of all the pre-fix requirements of the filter buffers.

In block 804, matrix switch 308 selectively removes one or more existing filter connections to adjacent filters. Matrix switch 308 then reconnects all of its pins to adjacent filters using a common buffer size between each of the pins, block 806. In block 808, matrix switch 308 negotiates to be the allocator for all of its pins (402, 404). If the matrix switch 308 cannot, for whatever reason, be the allocator for any of its input pins 402 minimal loss to performance is encountered, as the buffer associated with the input pin will still be compatible with any downstream filter (i.e., coupled to an output pin) and, thus, the buffer can still be passed to the downstream filter without requiring a memory copy operation. If, however, matrix switch 308 cannot be an allocator for one of its output pins 404, media content must then be transferred to at least the downstream filter associated with that output pin using a memory copy operation, block 810.

In block 812, once the matrix switch 308 has re-established its connection to adjacent filters, render engine 222 restores the connection in remaining filters using negotiated buffer requirements emanating from the matrix switch filter 308 buffer negotiations. Once the connections throughout the filter graph have been reconnected, the process continues with block 714 of Fig. 7.

1 In block 714 (Fig. 7), have re-established the connections between filters,
2 render engine 222 is ready to implement a user's instruction to execute the media
3 processing project.

4 5 **Example Operation and Implementation(s)**

6 The matrix switch described above is quite useful in that it allows multiple
7 inputs to be directed to multiple outputs at any one time. These input can compete
8 for a matrix switch output. The embodiments described below permit these
9 competing inputs to be organized so that the inputs smoothly flow through the
10 matrix switch to provide a desired output. And, while the inventive programming
11 techniques are described in connection with the matrix switch as such is employed
12 in the context of multi-media editing projects, it should be clearly understood that
13 application of the inventive programming techniques and structures should not be
14 so limited only to application in the field of multi-media editing projects or, for
15 that matter, multi-media applications or data streams. Accordingly, the principles
16 about to be discussed can be applied to other fields of endeavor in which multiple
17 inputs can be characterized as competing for a particular output during a common
18 time period.

19 In the multi-media example below, the primary output of the matrix switch
20 is a data stream that defines an editing project that has been created by a user.
21 Recall that this editing project can include multiple different sources that are
22 combined in any number of different ways, and the sources that make up a project
23 can comprise audio sources, video sources, or both. The organization of the inputs
24 and outputs of the matrix switch are made manageable, in the examples described
25 below, by a data structure that permits the matrix switch to be programmed.

1 Fig. 9 shows an overview of a process that takes a user-defined editing
2 project and renders from it a data structure that can be used to program the matrix
3 switch.

4 Specifically, a user-defined editing project is shown generally at 900.
5 Typically, when a user creates an editing project, they can select from a number of
6 different multimedia clips that they can then assemble into a unique presentation.
7 Each individual clip represents a *source* of digital data or a source stream (e.g.,
8 multimedia content). Projects can include one or more sources 902. In defining
9 their project, a user can operate on sources in different ways. For example, video
10 sources can have *transitions* 904 and *effects* 906 applied on them. A transition
11 object is a way to change between two or more sources. As discussed above, a
12 transition essentially receives as input, two or more streams, operates on them in
13 some way, and produces a single output stream. An exemplary transition can
14 comprise, for example, fading from one source to another. An effect object can
15 operate on a single source or on a composite of sources. An effect essentially
16 receives a single input stream, operates on it in some way, and produces a single
17 output stream. An exemplary effect can comprise a black-and-white effect in
18 which a video stream that is configured for presentation in color format is
19 rendered into a video stream that is configured for presentation in black and white
20 format. Unlike conventional effect filters, effect object 906 may well perform
21 multiple effect tasks. That is, in accordance with one implementation, an effect
22 object (e.g., 906) may actually perform multiple tasks on the received input
23 stream, wherein said tasks would require multiple effect filters in a conventional
24 filter graph system.
25

1 An exemplary user interface 908 is shown and represents what a user might
2 see when they produce a multimedia project with software executing on a
3 computer. In this example, the user has selected three sources A, B, and C, and
4 has assembled the sources into a project timeline. The project timeline defines
5 when the individual sources are to be rendered, as well as when any transitions
6 and/or effects are to occur.

7 In the discussion that follows, the notion of a *track* is introduced. A track
8 can contain one or more sources or source clips. If a track contains more than one
9 source clip, the source clips cannot overlap. If source clips are to overlap (e.g.
10 fading from one source to another, or having one source obscure another), then
11 multiple tracks are used. A track can thus logically represent a layer on which
12 sequential video is produced. User interface 908 illustrates a project that utilizes
13 three tracks, each of which contains a different source. In this particular project
14 source A will show for a period of time. At a defined time in the presentation,
15 source A is obscured by source B. At some later time, source B transitions to
16 source C.

17 In accordance with the described embodiment, the user-defined editing
18 project 900 is translated into a data structure 910 that represents the project. In the
19 illustrated and described example, this data structure 910 comprises a tree
20 structure. It is to be understood, however, that other data structures could be used.
21 The use of tree structures to represent editing projects is well-known and is not
22 described here in any additional detail. Once the data structure 910 is defined, it is
23 processed to provide a data structure 912 that is utilized to program the matrix
24 switch. In the illustrated and described embodiment, data structure 912 comprises
25 a grid from which the matrix switch can be programmed. It is to be understood

1 and appreciated that other data structures and techniques could, however, be used
2 to program the matrix switch without departing from the spirit and scope of the
3 claimed subject matter.

4 The processing that takes place to define data structures 910 and 912 can
5 take place using any suitable hardware, software, firmware, or combination
6 thereof. In the examples set forth below, the processing takes place utilizing
7 software in the form of a video editing software package that is executable on a
8 general purpose computer.

10 Example Project

11 For purposes of explanation, consider Fig. 10 which shows project 908
12 from Fig. 9 in a little additional detail. Here, a time line containing numbers 0-16
13 is provided adjacent the project to indicate when particular sources are to be seen
14 and when transitions and effects (when present) are to occur. In the examples in
15 this document, the following convention exists with respect to projects, such as
16 project 908. A priority exists for video portions of the project such that as one
17 proceeds from top to bottom, the priority increases. Thus, in the Fig. 10 example,
18 source A has the lowest priority followed by source B and source C. Thus, if there
19 is an overlap between higher and lower priority sources, the higher priority source
20 will prevail. For example, source B will obscure source A from between $t = 4-8$.

21 In this example, the following can be ascertained from the project 908 and
22 time line: from time $t=0-4$ source A should be routed to the matrix switch's
23 primary output; from $t=4-12$ source B should be routed to the matrix switch's
24 primary output; from $t=12-14$ there should be a transition between source B and
25 source C which should be routed to the matrix switch's primary output; and from

$t=14-16$ source C should be routed to the matrix switch's primary output. Thus, relative to the matrix switch, each of the sources and the transition can be characterized by where it is to be routed at any given time. Consider, for example, the table just below:

Object	Routing for a given time
C	$t = 0-12$ (nowhere); $t = 12-14$ (transition); $t = 14-16$ (primary output)
B	$t = 0-4$ (nowhere); $t = 4-12$ (primary output); $t = 12-14$ (transition); $t = 14-16$ (nowhere)
A	$t = 0-4$ (primary output); $t = 4-16$ (nowhere)
Transition	$t = 0-12$ (nowhere); $t = 12-14$ (primary output); $t = 14-16$ (nowhere)

Fig. 11 shows an exemplary matrix switch 1100 that can be utilized in the presentation of the user's project. Matrix switch 1100 comprises multiple inputs and multiple outputs. Recall that a characteristic of the matrix switch 1100 is that any of the inputs can be routed to any of the outputs at any given time. A transition element 1102 is provided and represents the transition that is to occur between sources B and C. Notice that the matrix switch includes four inputs numbered 0-3 and three outputs numbered 0-2. Inputs 0-2 correspond respectively to sources A-C, while input 3 corresponds to the output of the transition element 1102. Output 0 corresponds to the switch's primary output, while outputs 1 and 2 are routed to the transition element 1102.

The information that is contained in the table above is the information that is utilized to program the matrix switch. The discussion presented below describes

1 but one implementation in which the information contained in the above table can
2 be derived from the user's project time line.

3 Recall that as a user edits or creates a project, software that comprises a part
4 of their editing software builds a data structure that represents the project. In the
5 Fig. 9 overview, this was data structure 910. In addition to building the data
6 structure that represents the editing project, the software also builds and configures
7 a matrix switch that is to be used to define the output stream that embodies the
8 project. Building and configuring the matrix switch can include building the
9 appropriate graphs (e.g., a collection of software objects, or filters) that are
10 associated with each of the sources and associating those graphs with the correct
11 inputs of the matrix switch. In addition, building and configuring the matrix
12 switch can also include obtaining and incorporating additional appropriate filters
13 with the matrix switch, e.g. filters for transitions, effects, and mixing (for audio
14 streams). This will become more apparent below.

15 Fig. 12 shows a graphic representation of an exemplary data structure 1200
16 that represents the project of Fig. 10. Here, the data structure comprises a
17 traditional hierarchical tree structure. Any suitable data structure can, however, be
18 utilized. The top node 1202 constitutes a *group* node. A *group* encapsulates a type
19 of media. For example, in the present example the media type comprises video.
20 Another media type is audio. The group node can have child nodes that are either
21 tracks or composites. In the present example, three track nodes 1204, 1206, and
22 1208 are shown. Recall that each track can have one or more sources. If a track
23 comprises more than one source, the sources cannot overlap. Here, all of the
24 sources (A, B, and C) overlap. Hence, three different tracks are utilized for the
25 sources. In terms of priority, the lowest priority source is placed into the tree

1 furthest from the left at 1204a. The other sources are similarly placed. Notice that
2 source C (1208a) has a transition 1210 associated with it. A transition object, in
3 this example, defines a two-input/one output operation. When applied to a track
4 or a composition (discussed below in more detail), the transition object will
5 operate between the track to which it has been applied, and any objects that are
6 beneath it in priority and at the same level in the tree. A "tree level" has a
7 common depth within the tree and belongs to the same parent. Accordingly, in
8 this example, the transition 1210 will operate on a source to the left of the track on
9 which source C resides, and beneath it in priority, i.e. source B. If the transition is
10 applied to any object that has nothing beneath it in the tree, it will transition from
11 blackness (and/or silence if audio is included).

12 Once a data structure representing the project has been built, in this case a
13 hierarchical tree structure, a rendering engine processes the data structure to
14 provide another data structure that is utilized to program the matrix switch. In the
15 Fig. 9 example, this additional data structure is represented at 912. It will be
16 appreciated and understood that the nodes of tree 1200 can include so-called meta
17 information such as a name, ID, and a time value that represents when that
18 particular node's object desires to be routed to the output, e.g. node 1204a would
19 include an identifier for the node associating it with source A, as well as a time
20 value that indicates that source A desires to be routed to the output from time $t = 0$ -
21 8. This meta information is utilized to build the data structure that is, in turn,
22 utilized to program the matrix switch.

23 In the example about to be described below, a specific data structure in the
24 form of a grid is utilized. In addition, certain specifics are described with respect
25 to how the grid is processed so that the matrix switch can be programmed. It is to

1 be understood that the specific described approach is for exemplary purposes only
2 and is not intended to limit application of the claims. Rather, the specific approach
3 constitutes but one way of implementing broader conceptual notions embodied by
4 the inventive subject matter.

5 Figs. 13-18 represent a process through which the inventive grid is built. In
6 the grid about to be described, the x axis represents time, and the y axis represents
7 layers in terms of priority that go from lowest (at the top of the grid) to highest (at
8 the bottom of the grid). Every row in the grid represents the video layer.
9 Additionally, entries made within the grid represent output pins of the matrix
10 switch. This will become apparent below.

11 The way that the grid is built in this example is that the rendering engine
12 does a traversal operation on the tree 1200. In this particular example, the
13 traversal operation is known as a “depth-first, left-to-right” traversal. This
14 operation will layerize the nodes so that the leftmost track or source has the lowest
15 priority and so on. Doing the above-mentioned traversal on tree 1200 (Fig. 12),
16 the first node encountered is node 1204 which is associated with source A. This is
17 the lowest priority track or source. A first row is defined for the grid and is
18 associated with source A. After the first grid row is defined, a grid entry is made
19 and represents the time period for which source A desires to be routed to the
20 matrix switch’s primary output.

21 Fig. 13 shows the state of a grid 1300 after this first processing step.
22 Notice that from time $t = 0-8$, a “0” has been placed in the grid. The “0”
23 represents the output pin of the matrix switch—in this case the primary output.
24 Next, the traversal encounters node 1206 (Fig. 12) which is associated with source
25 B. A second row is thus defined for the grid and is associated with source B.

After the second grid row is defined, a grid entry is made and represents the time period for which source B desires to be routed to the matrix switch's primary output.

Fig. 14 shows the state of grid 1300 after this second processing step. Notice that from time $t = 4-14$, a "0" has been placed in the grid. Notice at this point that something interesting has occurred which will be resolved below. Each of the layers has a common period of time (i.e. $t = 4-8$) for which it desires to be routed to the matrix switch's primary output. However, because of the nature of the matrix switch, only one input can be routed to the primary output at a time. Next, the traversal encounters node 1208 (Fig. 12) which is associated with source C. In this particular processing example, a rule is defined that sources on tracks are processed before transitions on the tracks are processed because transitions operate on two objects that are beneath them. A third row is thus defined for the grid and is associated with source C. After the third row is defined, a grid entry is made and represents the time period for which source C desires to be routed to the matrix switch's primary output.

Fig. 15 shows the state of grid 1300 after this third processing step. Notice that from time $t = 12-16$, a "0" has been placed in the grid. Next, the traversal encounters node 1210 (Fig. 12) which corresponds to the transition. Thus, a fourth row is defined in the grid and is associated with the transition. After the fourth row is defined, a grid entry is made and represents the time period for which the transition desires to be routed to the matrix switch's primary output.

Fig. 16 shows the state of grid 1300 after this fourth processing step. Notice that from time $t = 12-14$, a "0" has been placed in the grid for the transition entry. The transition is a special grid entry. Recall that the transition is

1 programmed to operate on two inputs and provide a single output. Accordingly,
2 starting at the transition entry in the grid and working backward, each of the
3 entries corresponding to the same tree level are examined to ascertain whether
4 they contain entries that indicate that they want to be routed to the output during
5 the same time that the transition is to be routed to the output. If grid entries are
6 found that conflict with the transition's grid entry, the conflicting grid entry is
7 changed to a value that corresponds to an output pin that serves as an input to the
8 transition element 1102 (Fig. 11). This is essentially a redirection operation. In
9 the illustrated grid example, the transition first finds the level that corresponds to
10 source C. This level conflicts with the transition's grid entry for the time period t
11 $= 12-14$. Thus, for this time period, the grid entry for level C is changed to a
12 switch output that corresponds to an input for the transition element. In this
13 example, a "2" is placed in the grid to signify that for this given time period, this
14 input is routed to output pin 2. Similarly, continuing up the grid, the next level
15 that conflicts with the transition's grid entry is the level that corresponds to source
16 B. Thus, for the conflicting time period, the grid entry for level B is changed to a
17 switch output that corresponds to an input for the transition element. In this
18 example, a "1" is placed in the grid to signify that for this given time period, this
19 input is routed to output pin 1 of the matrix switch.

20 Fig. 17 shows the state of the grid at this point in the processing. Next, a
21 pruning function is implemented which removes any other lower priority entry
22 that is contending for the output with a higher priority entry. In the example, the
23 portion of A from $t=4-8$ gets removed because the higher priority B wants the
24 output for that time.
25

1 Fig. 18 shows the grid with a cross-hatched area that signifies that portion
2 of A's grid entry that has been removed.

3 At this point, the grid is in a state in which it can be used to program the
4 matrix switch. The left side entries -- A, B, C, and TRANS represent input pin
5 numbers 0, 1, 2, and 3 (as shown) respectively, on the matrix switch shown in Fig.
6 11. The output pin numbers of the matrix switch are designated at 0, 1, and 2 both
7 on the switch in Fig. 11 and within the grid in Fig. 18. As one proceeds through
8 the grid, starting with source A, the programming of the matrix switch can be
9 ascertained as follows: A is routed to output pin 0 of the matrix switch (the
10 primary output) from $t = 0-4$. From $t = 4-16$, A is not routed to any output pins.
11 From $t = 0-4$, B is not routed to any of the output pins of the matrix switch. From t
12 $= 4-12$, B is routed to the primary output pin 0 of the matrix switch. From $t = 12-$
13 14 , B is routed to output pin 1 of the matrix switch. Output pin 1 of the matrix
14 switch corresponds to one of the input pins for the transition element 1102 (Fig.
15 11). From $t = 14-16$, B is not routed to any of the output pins of the matrix switch.
16 From $t = 0-12$, C is not routed to any of the output pins of the matrix switch. From
17 $t = 12-14$, C is routed to output pin 2 of the matrix switch. Output pin 2 of the
18 matrix switch corresponds to one of the input pins for the transition element 302
19 (Fig. 3). From $t = 12-14$ the transition element (input pin 3) is routed to output pin
20 0. From $t = 14-16$, C is routed to output pin 0 of the matrix switch.

21 As alluded to above, one of the innovative aspects of the matrix switch 308
22 is its ability to seek to any point in a source, without having to process the
23 intervening content serially through the filter. Rather, matrix switch 308 identifies
24 an appropriate transition point and dumps at least a subset of the intervening
25 content, and continues processing from the sought point in the content.

1 The ability of the matrix switch 308 to seek to any point in the media
2 content gives rise to certain performance enhancement heretofore unavailable in
3 computer implemented media processing systems. For example, generation of a
4 filter graph by render engine 222 may take into account certain performance
5 characteristics of the media processing system which will execute the user-defined
6 media processing project. In accordance with this example implementation,
7 render engine 222 may access and analyze the system registry of the operating
8 system, for example, to ascertain the performance characteristics of hardware
9 and/or software elements of the computing system implementing the media
10 processing system, and adjust the filter graph construction to improve the
11 perceived performance of the media processing system by the user. Nonetheless,
12 there will always be a chance that a particular instance of a filter graph will not be
13 able to process the media stream fast enough to provide the desired output at the
14 desired time, i.e., processing of the media stream bogs down leading to delays at
15 the rendering filter. In such a case, matrix switch 308 will recognize that it is not
16 receiving media content at the appropriate project time, and may skip certain
17 sections of the project in an effort to "catch-up" and continue the remainder of the
18 project in real time. According to one implementation, when matrix switch 308
19 detects such a lag in processing, it will analyze the degree of the lag and issue a
20 seek command to the source (through the source processing chain) to a future
21 point in the project, where processing continues without processing any further
22 content prior to the seeked point.

23 Thus, for the editing project depicted in Fig. 10, the processing described
24 above first builds a data structure (i.e. data structure 1200 in Fig. 12) that
25 represents the project in hierarchical space, and then uses this data structure to

1 define or create another data structure that can be utilized to program the matrix
2 switch.

3 Fig. 19 is a flow diagram that describes steps in a method in accordance
4 with the described embodiment. The method can be implemented in any suitable
5 hardware, software, firmware, or combination thereof. In the illustrated and
6 described embodiment, the method is implemented in software.

7 Step 1900 provides a matrix switch. An exemplary matrix switch is
8 described above. Step 1902 defines a first data structure that represents the editing
9 project. Any suitable data structure can be used, as will be apparent to those of
10 skill in the art. In the illustrated and described embodiment, the data structure
11 comprises a hierarchical tree structure having nodes that can represent tracks
12 (having one or more sources), composites, transitions and effects. Step 1904
13 processes the first data structure to provide a second data structure that is
14 configured to program the matrix switch. Any suitable data structure can be
15 utilized to implement the second data structure. In the illustrated and described
16 embodiment, a grid structure is utilized. Exemplary processing techniques for
17 processing the first data structure to provide the second data structure are
18 described above. Step 1906 then uses the second data structure to program the
19 matrix switch.

20 21 Example Project with a Transition and an Effect

22 Consider project 2000 depicted in Fig. 20. In this project there are three
23 tracks, each of which contains a source, i.e. source A, B and C. This project
24 includes an effect applied on source B and a transition between sources B and C.
25 The times are indicated as shown.

As the user creates their project, a data structure representing the project is built. Fig. 21 shows an exemplary data structure in the form of a hierarchical tree 2100 that represents project 2000. There, the data structure includes three tracks, each of which contains one of the sources. The sources are arranged in the tree structure in the order of their priority, starting with the lowest priority source on the left and proceeding to the right. There is an effect (i.e. "Fx") that is attached to or otherwise associated with source B. Additionally, there is a transition attached to or otherwise associated with source C.

In building the grid for project 2000, the following rule is employed for effects. An effect, in this example, is a one-input/one-output object that is applied to one object—in this case source B. When the effect is inserted into the grid, it looks for any one object beneath it in priority that has a desire to be routed to the primary output of the matrix switch at the same time. When it finds a suitable object, it redirects that object's output from the matrix switch's primary output to an output associated with the effect.

As an example, consider Fig. 22 and the grid 2200. At this point in the processing of tree 2100, the rendering engine has incorporated entries in the grid corresponding to sources A, B and the effect. It has done so by traversing the tree 2100 in the above-described way. In this example, the effect has already looked for an object beneath it in priority that is competing for the primary output of the matrix switch. It found an entry for source B and then redirected B's grid entry to a matrix switch output pin that corresponds to the effect—here output pin 1.

As the render engine 222 completes its traversal of tree 2100, it completes the grid. Fig. 23 shows a completed grid 2200. Processing of the grid after that which is indicated in Fig. 22 takes place substantially as described above with

respect to the first example. Summarizing, this processing though: after the effect is entered into the grid and processed as described above, the traversal of tree 2100 next encounters the node associated with source C. Thus, a row is added in the grid for source C and an entry is made to indicate that source C desires the output from $t = 12-16$. Next, the tree traversal encounters the node associated with the transition. Accordingly, a row is added to the grid for the transition and a grid entry is made to indicate that the transition desires the output from $t = 12-14$. Now, as described above, the grid is examined to find two entries, lower in priority than the transition and located at the same tree level as the transition, that compete for the primary output of the matrix switch. Here, those entries correspond to the grid entries for the effect and source C that occur from $t = 12-14$. These grid entries are thus redirected to output pins of the matrix switch 308 that correspond to the transition—here pins 2 and 3 as indicated. Next, the grid is pruned which, in this example, removes a portion of the grid entry corresponding to source A for $t = 4-8$ because of a conflict with the higher-priority entry for source B.

Fig. 24 shows the resultant matrix switch that has been built and configured as the grid was being processed above. At this point, the grid can be used to program the matrix switch. From the grid picture, it is very easy to see how the matrix switch 308 is going to be programmed. Source A will be routed to the matrix switch's primary output (pin 0) from $t = 0-4$; source B will be redirected to output pin 1 (effect) from $t = 4-14$ and the effect on B will be routed to the output pin 0 from $t = 4-12$. From $t = 12-14$, the effect and source C will be routed to output pins corresponding to the transition (pins 2 and 3) and, accordingly, during this time the transition (input pin 4) will be routed to the primary output (output

pin 0) of the matrix switch. From $t = 14-16$, source C will be routed to the primary output of the matrix switch.

It will be appreciated that as the software, in this case the render engine 222, traverses the tree structure that represents a project, it also builds the appropriate graphs and adds the appropriate filters and graphs to the matrix switch. Thus, for example, as the render engine 222 encounters a tree node associated with source A, in addition to adding an entry to the appropriate grid, the software builds the appropriate graphs (i.e. collection of linked filters), and associates those filters with an input of the matrix switch. Similarly, when the render engine 222 encounters an effect node in the tree, the software obtains an effect object or filter and associates it with the appropriate output of the matrix switch. Thus, in the above examples, traversal of the tree structure representing the project also enables the software to construct the appropriate graphs and obtain the appropriate objects and associate those items with the appropriate inputs/outputs of the matrix switch 308. Upon completion of the tree traversal and processing of the grid, an appropriate matrix switch has been constructed, and the programming (i.e. timing) of inputs to outputs for the matrix switch has been completed.

Treatment of “blanks” in a Project

There may be instances in a project when a user leaves a blank in the project time line. During this blank period, no video or audio is scheduled for play.

Fig. 25 shows a project that has such a blank incorporated therein. If there is such a blank left in a project, the software is configured to obtain a “black” source and associate the source with the matrix switch at the appropriate input pin.

1 The grid is then configured when it is built to route the black source to the output
2 at the appropriate times and fade from the black (and silent) source to the next
3 source at the appropriate times. The black source can also be used if there is a
4 transition placed on a source for which there is no additional source from which to
5 transition.

6 7 Audio Mixing

8 In the examples discussed above, sources comprising video streams were
9 discussed. In those examples, at any one time, only two video streams were
10 combined into one video stream. However, each project can, and usually does
11 contain an audio component. Alternately, a project can contain only an audio
12 component. The audio component can typically comprise a number of different
13 audio streams that are combined. The discussion below sets forth but one way of
14 processing and combining audio streams.

15 In the illustrated example, there is no limit on the number of audio streams
16 that can be combined at any one time.

17 Suppose, for example, there is an audio project that comprises 5 tracks, A-
18 E. Fig. 26 shows an exemplary project. The shaded portions of each track
19 represent the time during which the track is not playing. So, for example, at $t=0-4$,
20 tracks B, D, and E are mixed together and will play. From $t = 4-10$, tracks A-E are
21 mixed together and will play, and the like.

22 Fig. 27 shows the grid for this project at 2700. Since we are dealing with
23 this composition now, all of the effects and transitions including the audio mixing
24 are only allowed to affect things in this composition. Thus, there is the concept of
25 a boundary 2702 that prevents any actions or operations in this composition from

1 affecting any other grid entries. Note that there are other entries in the grid and
2 that the presently-illustrated entries represent only those portions of the project
3 that relate to the audio mixing function.

4 Grid 2700 is essentially set up in a manner similar to that described above
5 with respect to the video projects. That is, for each track, a row is added to the
6 grid and a grid entry is made for the time period during which the source on that
7 track desires to be routed to the primary output of the matrix switch. In the
8 present example, grid entries are made for sources A-E. Next, in the same way
9 that a transition or effect was allocated a row in the grid, a "mix" element is
10 allocated a row in the grid as shown and a grid entry is made to indicate that the
11 mix element desires to be routed to the primary output of the matrix switch for a
12 period of time during which two or more sources compete for the matrix switch's
13 primary output. Note that in this embodiment, allocation of a grid row for the mix
14 element can be implied. Specifically, whereas in the case of a video project,
15 overlapping sources simply result in playing the higher priority source (unless the
16 user defines a transition between them), in the audio realm, overlapping sources
17 are treated as an implicit request to mix them. Thus, the mix element is allocated a
18 grid row any time there are two or more overlapping sources.

19 Once the mix element is allocated into the grid, the grid is processed to
20 redirect any conflicting source entries to matrix switch output pins that correspond
21 to the mix element. In the above case, redirection of the grid entries starts with pin
22 3 and proceeds through to pin 7. The corresponding matrix switch is shown in
23 Fig. 28. Notice that all of the sources are now redirected through the mix element
24 which is a multi-input/one output element. The mix element's output is fed back
25 around and becomes input pin 15 of the matrix switch. All of the programming of

1 the matrix switch is now reflected in the grid 2700. Specifically, for the indicated
2 time period in the grid, each of the sources is routed to the mix element which, in
3 turn, mixes the appropriate audio streams and presents them to the primary output
4 pin 0 of the matrix switch.

5 6 Compositions

7 There are situations that can arise when building an editing project where it
8 would be desirable to apply an effect or a transition on just a subset of a particular
9 project or track. Yet, there is no practicable way to incorporate the desired effect
10 or transition. In the past, attempts to provide added flexibility for editing projects
11 have been made in the form of so called "bounce tracks", as will be appreciated
12 and understood by those of skill in the art. The use of bounce tracks essentially
13 involves processing various video layers (i.e. tracks), writing or moving the
14 processed layers or tracks to another location, and retrieving the processed layers
15 when later needed for additional processing with other layers or tracks. This type
16 of processing can be slow and inefficient.

17 To provide added flexibility and efficiency for multi-media editing projects,
18 the notion of a *composite* or *composition* is introduced. A composite or
19 composition can be considered as a representation of an editing project as a single
20 track. Recall that editing projects can have one or more tracks, and each track can
21 be associated with one or more sources that can have effects applied on them or
22 transitions between them. In addition, compositions can be nested inside one
23 another.

24 25 Example Project with Composite

Consider, for example, Fig. 29 which illustrates an exemplary project 2900 having a composition 2902. In this example, composition 2902 comprises sources B and C and a transition between B and C that occurs between $t = 12-14$. This composition is treated as an individual track or layer. Project 2900 also includes a source A, and a transition between source A and composition 2902 at $t = 4-8$. It will be appreciated that compositions can be much more complicated than the illustrated composition, which is provided for exemplary purposes only. Compositions are useful because they allow the grouping of a particular set of operations on one or more tracks. The operation set is performed on the grouping, and does not affect tracks that are not within the grouping. To draw an analogy, a composition is similar in principle to a mathematical parenthesis. Those operations that appear within the parenthesis are carried out in conjunction with those operations that are intended to operate of the subject matter of the parenthesis. The operations within the parenthesis do not affect tracks that do not appear within the parenthesis.

In accordance with the processing that is described above in connection with Fig. 19, a first data structure is defined that represents the editing project. Fig. 30 shows an exemplary data structure 3000 in the form of a hierarchical tree structure. In this example, group node 3002 includes two children—track node 3004 and composite node 3006. Track node 3004 is associated with source A. Composite node 3006 includes two children—track nodes 3008 and 3010 that are respectively associated with sources B (3008a) and C (3010a). A transition T2 (3012) is applied on source C and a transition T1 (3014) is applied on composition 3006.

Next, data structure 3000 is processed to provide a second data structure that is configured to program the matrix switch. Note that as the data structure is being programmed, a matrix switch is being built and configured at the same time. In this example, the second data structure comprises a grid structure that is assembled in much the same way as was described above. There are, however, some differences and, for purposes of understanding, the complete evolution of the grid structure is described here. In the discussion that follows, the completed matrix switch is shown in Fig. 38.

When the rendering engine initiates the depth-first, left-to-right traversal of data structure 3000, the first node it encounters is track node 3004 which is associated with source A. Thus, a first row of the grid is defined and a grid entry is made that represents the time period for which source A desires to be routed to the matrix switch's primary output pin.

Fig. 31 shows the state of a grid 3100 after this first processing step. Next the traversal of data structure 3000 encounters the composite node 3006. The composite node is associated with two tracks—track 3008 and track 3010. Track 3008 is associated with source B. Accordingly, a second row of the grid is defined and a grid entry is made that represents the time period for which source B desires to be routed to the matrix switch's primary output pin. Additionally, since B is a member of a composition, meta-information is contained in the grid that indicates that this grid row defines one boundary of the composition. This meta-information is graphically depicted with a bracket that appears to the left of the grid row.

Fig. 32 shows the state of grid 3100 after this processing step. Next, the traversal of data structure 3000 encounters node 3010 which is associated with

1 source C. Thus, a third row of the grid is added and a grid entry is made that
2 represents the time period for which source C desires to be routed to the matrix
3 switch's primary output pin.

4 Fig. 33 shows the state of grid 3100 after this processing step. Notice that
5 the bracket designating the composition now encompasses the grid row associated
6 with source C. The traversal next encounters node 3012 which is the node
7 associated with the *second* transition T2. Thus, as in the above example, a grid
8 row is added for the transition and a grid entry is made that represents the time
9 period for which the transition desires to be routed to the matrix switch's primary
10 output pin.

11 Fig. 34 shows the state of grid 3100 after this processing step. Notice that
12 the bracket designating the composition is now completed and encompasses grid
13 row entries that correspond to sources B and C and the transition between them.
14 Recall from the examples above that a transition, in this example, is programmed
15 to operate on two inputs and provide a single output. In this instance, and because
16 the transition occurs within a composition, the transition is constrained by a rule
17 that does not allow it to operate on any elements outside of the composition.
18 Thus, starting at the transition entry and working backward through the grid,
19 entries at the same tree level and within the composition (as designated by the
20 bracket) are examined to ascertain whether they contain entries that indicate that
21 they want to be routed to the output during the same time that the transition is to
22 be routed to the output. Here, both of the entries for sources B and C have
23 portions that conflict with the transition's entry. Accordingly, those portions of
24 the grid entries for sources B and C are redirected or changed to correspond to
25

1 output pins that are associated with a transition element that corresponds to
2 transition T2.

3 Fig. 35 shows the state of grid 3100 after this processing step. The
4 traversal next encounters node 3014 which is the node that is associated with the
5 transition that occurs between source A and composition 2902 (Fig. 29).
6 Processing of this transition is similar to processing of the transition immediately
7 above except for the fact that the transition does not occur within the composition.
8 Because the transition occurs between the composition and another source, one of
9 the inputs for the transition will be the composition, and one of the inputs will be
10 source A (which is outside of the composition). Thus, a grid row is added for this
11 transition and a grid entry is made that represents the time period for which the
12 transition desires to be routed to the matrix switch's primary output pin.

13 Fig. 36 shows the state of grid 3100 after this processing step. At this point
14 then, the grid is examined for entries that conflict with the entry for transition T1.
15 One conflicting grid entry is found for the row that corresponds to source B (inside
16 the composition) and one that corresponds to source A (outside the composition).
17 Accordingly, those portions of the grid row that conflict with transition T1 are
18 changed or redirected to have values that are associated with output pins of the
19 matrix switch that are themselves associated with a transition element T1. In this
20 example, redirection causes an entry of "3" and "4" to be inserted as shown.

21 Fig. 37 shows the state of grid 3100 after this processing step. If necessary,
22 a pruning operation would further ensure that the grid has no competing entries for
23 the primary output of the matrix switch. The associated input pin numbers of the
24 matrix switch are shown to the left of grid 3100.
25

Fig. 38 shows a suitably configured matrix switch that has been build in accordance with the processing described above. Recall that, as data structure 3000 (Fig. 30) is processed by the rendering engine, a matrix switch is built and configured in parallel with the building and processing of the grid structure that is utilized to program the matrix switch. From the matrix switch and grid 3100 of Fig. 37, the programming of the switch can be easily ascertained.

Fig. 38a shows an exemplary data structure that represents a project that illustrates the usefulness of composites. In this example, the project can mathematically be represented as follows:

$$(Fx\text{-noisy } (A \text{ Tx-Blend } B)) \text{ Tx-Blend } C$$

Here, an effect (noisy) is applied to A blended with B, the result of which is applied to a blend with C. The composite in this example allows the grouping of the things beneath it so that the effect (noisy), when it is applied, is applied to everything that is beneath it. Notice that without the composite node, there is no node where an effect can be applied that will affect (A Tx-Blend B). Hence, in this example, operations that appear within the parenthesis are carried out on tracks that appear within the parenthesis. Those operations do not affect tracks that are not within the parenthesis.

Fig. 39 is a flow diagram that described steps in a method in accordance with one embodiment. The method can be implemented in any suitable hardware, software, firmware, or combination thereof. In the presently-described example, the method is implemented in software.

Step 3900 defines a multimedia editing project that includes at least one composite. The composite represents multiple tracks as a single track for purposes of the processing described just below. It is important to note that, in the processing described just below, and because of the use of composites, the extra processing that is required by bounce tracks is avoided (i.e. operating on two tracks, moving the operation result to another location, and retrieving the operation result when later needed). This reduces the processing time that is required to render a multi-media project. Step 3902 defines a first data structure that represents the editing project. Any suitable data structure can be utilized. In the present example, a data structure in the form of a hierarchical tree is utilized. An exemplary tree is shown in Fig. 30. Step 3904 processes the first data structure to provide a second data structure that is configured to program a matrix switch. In the illustrated example, the second data structure comprises a grid structure. Exemplary processing is described in the context of Figs. 30-37. Step 3906 then programs the matrix switch using the second data structure.

Source Combining

Having introduced the various architectural and implementation elements of the present invention, above, attention is now drawn to Figs. 40-43, wherein another aspect of the illustrated embodiment is presented. As introduced above, the opening and processing of media represents consumption of memory and processing resources. Thus, performance improvements may be achieved by reducing the number of times a source is accessed. Thus, a method is presented in accordance with one aspect of the present invention that serves to reduce the number of times a source is accessed, e.g., a method of source combining. It is to

1 be appreciated, however, that the following is but one example implementation of
2 the broader inventive concept of reducing the number of times a source need be
3 accessed during execution of a development project. Alternative methods of
4 source combining of greater or lesser complexity may well be used within the
5 spirit and scope of the present invention. Indeed, such alternative methods are
6 anticipated within the scope of the present invention.

7 **Fig. 41** illustrates an example method of generating a filter graph, in
8 accordance with one aspect of the present invention. As shown, method 4000
9 begins with block 4002, wherein render engine 222 receives an indication to
10 generate a development project. According to one implementation, as discussed
11 above, render engine 222 receives the indication from a higher-level application
12 216, e.g., media processing application 224, to assist a user in generating a
13 processing project (e.g., a media processing project).

14 In block 4004, render engine 222 identifies the number and nature of the
15 media sources within the user-defined processing project, in preparation for
16 generating a filter graph representation of the processing project. As introduced
17 above, for each of the identified sources, render engine 222 determines the
18 necessary transform filters 306 required to pre-process the source (i.e., source
19 chain), preparing the processing chain for presentation to the matrix switch filter
20 308 and one or more transition/effect filters 306. Unlike conventional
21 implementations, which would proceed to generate the entire filter graph in
22 preparation for execution of the processing project, render engine 222 generates a
23 list of sources and when they are required in the filter graph. According to one
24 implementation, the list is referred to as a reuse list, and is maintained within
25

1 render engine 222. An example of a data structure comprising a reuse list is
2 presented with reference to Fig. 41.

3 Turning briefly to Fig. 41, a graphical illustration of an example data
4 structure comprising a source reuse list is presented. As shown, the reuse list 4100
5 is comprised of a number of information fields, e.g., 4102-4110 which detail, in
6 part, the relationship between clips in a track. More particularly, the reuse list
7 4100 is shown comprising a track identification field 4102, a source identification
8 field 4104, a project time field 4106 and a source time field 4108.

9 Upon identifying a project source and the associated filters required for pre-
10 processing the source (i.e., the source chain), render engine 222 assigns each track
11 an identifier which uniquely identifies the source track within the context of the
12 filter graph. In this regard, reuse list 4100 includes a field 4102 which maintains a
13 list of tracks utilized in the associated project. In accordance with the illustrated
14 example paradigm of the media processing system, the track identifier is utilized
15 to represent a media clip from a given source.

16 The source identifier field 4104 contains information denoting the project
17 source associated with a particular track identifier. In this regard, the source
18 identifier field 4104 may well contain a file name, a file handle, or any other
19 suitable source identifier.

20 The project time field 4106 denotes at what point during project execution
21 the media clip is required. The source time field 4108 denotes what portion of the
22 source file is required to support execution of the processing project. It should be
23 appreciated that a user may well utilize the whole source file or any part thereof,
24 as defined by the processing project.
25

1 In accordance with the illustrated example implementation of Fig. 41, two
2 tracks are depicted 4110 and 4112. As shown, each of the tracks represent media
3 from a common source (e.g., source ID 4213) and, the source media clips are
4 adjacent to one another in the project (e.g., project time 4106) as well as within the
5 source file (e.g., source time 4108). As will be developed more fully below,
6 source clips may well be combined in certain situations into a single clip, as is
7 represented by track 4114 in Fig. 41. It is to be appreciated that, although depicted
8 as a two-dimensional data structure, reuse lists of greater or lesser complexity may
9 well be substituted without deviating from the spirit and scope of the present
10 invention.

11 Returning to Fig 40 and, in particular, block 4006, render engine 222
12 reduces the number of source accesses where possible, in accordance with one
13 aspect of the present invention. More particularly, render engine 222 analyzes the
14 reuse list 4100 to identify opportunities to reduce the number of source accesses
15 by combining source clips which meet certain criteria. According to one
16 implementation, the criteria used by render engine 222 include one or more of (1)
17 the source clips must occur next to one another in the project, (2) the clips appear
18 next to one another in the source, and (3) the clips must share a common pre-
19 processing source chain (i.e., must require the same pre-processing). If this
20 criteria is met, render engine 222 may combine the clips into a single clip. More
21 specifically, render engine 222 modifies the reuse list 4100 (Fig. 41) to replace the
22 multiple source accesses (4110, 4112) with a single source access 4114
23 representing both source accesses as a single access. It is to be appreciated that
24 removing a source access improves filter graph performance and, accordingly, the
25 perceived performance of the development system by the user.

1 In block 4008, once render engine 222 has reduced the number of source
2 file accesses (block 4006), render engine 222 dynamically generates and manages
3 the filter graph to support execution of the development project. In accordance
4 with one aspect of the present invention, render engine 222 invokes only those
5 source chains associated with sources that are necessary to support the current
6 and/or impending execution of the filter graph. It is to be appreciated that by not
7 opening each of the chains of a processing project, render engine 222 reduces the
8 amount of memory required to build the filter graph, thereby reducing the amount
9 of memory required to complete execution of the project.

10 Turning to **Fig. 42**, an example method for source combining is presented,
11 in accordance with one aspect of the present invention. As shown, the method
12 begins with block 4202, wherein render engine 222 identifies adjacent clips from a
13 common source, i.e., project aligned clips. As introduced above, render engine
14 222 analyzes the reuse list 4100 to identify all of the clips associated with a
15 particular track. For each of the adjacent source clips within a track, render engine
16 222 determines whether the clips are adjacent to one another with respect to their
17 source time, block 4204. That is, identifying that the source clips are adjacent to
18 one another by project time (e.g., occurring within the same track), render engine
19 222 next determines whether the clips are adjacent to one another in the source
20 file, i.e., whether the clips are source aligned. According to one implementation, if
21 the source clips are project and source time aligned, render engine 222 determines
22 whether the source clips share a common preprocessing source chain (i.e., the
23 require the same pre-processing).

24 If the source clips are source aligned, render engine 222 next determines
25 whether the clips require unique pre-processing (e.g., decoding, frame rate

conversion, sizing, etc.), block 4206. If unique pre-processing is required (block 4206), or the adjacent project clips are not source aligned (block 4204), the source clips will require an independent source chain of filters and, thus cannot be source combined. Accordingly, the source clips are accessed and processed independently, block 4208.

If, in block 4206, render engine 222 determines that the clips are source aligned and each share common pre-processing requirements, render engine 222 combines the adjacent source clips into a single clip and updates the reuse list 4100 accordingly, block 4210 (e.g., clip 4114 of Fig. 41). The combined source clip is representative of each of the otherwise individual clips, while requiring a single source processing chain and, thus, a single source access. In block 4212, render engine 222 determines whether all of the clips of the reuse list has been analyzed and, if not, the process continues in an iterative manner with block 4202 until the entire reuse list has been analyzed and appropriate source clips combined.

Fig. 43 graphically illustrates the source combining aspect of the present invention. As shown, Fig. 43 illustrates a project 4300 of two tracks of clips (e.g., clips 4302-4318, and 4320-4324) separated by a transition 4318. In accordance with the source combining aspect of the present invention, introduced above, render engine 222 analyzes a reuse list 4100 representation of project 4300 to identify source clips which may be combined. As introduced above, render engine 222 will combine source clips which are project and source time aligned, and which do not require unique pre-processing.

In the illustrated example of Fig. 43, clips 4302 and 4304 are project aligned and source aligned. If such clips do not require independent pre-processing, they are combined into a single clip 4326 by render engine 222. Note

1 that although clips 4304 and 4308 are project aligned, they are not source aligned
2 (e.g., the media end time (9) of clip 4304 does not abut the media start time (10) of
3 clip 4308), i.e., there is a gap of one elemental unit (e.g., a second of time). The
4 process of source combining is performed for other clips in the development
5 project, reducing the total number of source clips from eleven to six in the
6 illustrated example of Fig. 43. Thus, it is to be appreciated that the source
7 combining aspect of the present invention represents another feature which
8 reduces the computational and memory requirements necessary to support even
9 the most complex development projects.

10 11 **Source Filter Reuse**

12 As introduced above, conventional media processing systems may generate
13 an individual thread each time content was required from a source, even if the
14 source had been accessed earlier. This redundant loading/unloading of a source is
15 computationally expensive, and consumes precious memory resources. Extending
16 the concept of source combining introduced above, a filter and related methods for
17 sharing a common source and source filter among multiple processing threads will
18 now be introduced, in accordance with yet another aspect of the present invention.

19 Turning to **Fig. 44**, a block diagram of an example filter graph 4400 is
20 presented incorporating a segment filter 4406 which, as will be shown,
21 dynamically couples a source filter to one or more processing chain in accordance
22 with the teachings of the present invention. In accordance with the illustrated
23 example of Fig. 44, filter graph 4400 is depicted comprising a source 4402, one or
24 more pre-processing transform filters 4404, a segment filter 4406 and one or more
25

1 pre-processing transform filter(s) 4408A-N, each coupled to a matrix switch 308
2 and rendering filter(s) 4410, 4412, respectively.

3 As used herein, segment filter 4406 is designed to sit between a source
4 filter and matrix switch 308 to provide multiple processing chains with source
5 content from a single source, where it is impossible to combine the source clips (as
6 introduced above). Render engine 222 invokes an instance of the segment filter
7 4406 after the greatest common pre-processing filter 4404 for each of the chains.
8 That is, each of the processing chains may require the source content in a different
9 format (e.g., size, frame rate, decode format, etc.). To the extent that the chains
10 share common pre-processing attributes, those filter(s) (4404) are placed before
11 the segment filter 4406 where practicable. In many instances, none of the chains
12 share common pre-processing and the pre-processing filter(s) merely comprise the
13 source filter.

14 The segment filter 4406 acts as a controller, or throttle for the source,
15 instructing the source filter to deliver content from source 4402 at select times.
16 According to one implementation, the segment filter 4406 is, in turn, controlled by
17 the render engine 222 and/or the matrix switch filter 308 to provide select content
18 at select times on select inputs of the matrix switch filter 308. According to one
19 implementation, the segment filter 4406 issues a "seek" command to the source
20 filter to request particular content from the source. The source filter then delivers
21 the requested content through the segment filter 4406 and appropriate pre-
22 processing filter(s) 4408A-N to deliver the desired content to the requesting matrix
23 switch 308 to support processing of the development project.

24 As introduced above, render engine 222 is responsive to higher-level user
25 interfaces, e.g., applications 224. In this regard, it is possible that the filter graph

will receive user-commands while the filter graph is executing the development project. In accordance with the media processing system paradigm, for example, it is foreseeable that a user-invoked seek will be received by the filter graph during execution of the development project. Such user defined commands are typically serialized with commands issued by filters within the filter graph during the normal course of execution. In accordance with the illustrated example implementation, where matrix switch 308 "throttles" execution of the filter graph, matrix switch 308 issues a seek command of its own to the source filter, requesting the information desired by the user. According to an alternate embodiment, seeks received from a higher-level application (and, therefore, representative of a user command) are afforded a higher priority within the filter graph. In such an implementation, all segment filters 4406 residing within the filter graph are also notified of such high-priority seeks, so that they can identify what content they will be required to provide next and, therefore, issue a revised seek command of their own.

The remaining pre-processing transform filter(s) 4408A-N, matrix switch filter(s) 308 and rendering filter(s) 4410 each function as described above.

Turning now to Fig. 45, an example method for generating a filter graph is presented incorporating the teachings of the present invention. More particularly, the method of Fig. 45 is similar to the method of Fig. 42 wherein render engine 222 attempted source combining of source clips, which were not project and source time aligned, or which required unique pre-processing of some sort. In Fig. 42, however, if the source clips were not source time aligned (4204) and/or the clips required separate pre-processing (block 4206), each clip was assigned to a

1 separate processing chain. In Fig. 45, however, this problem is resolved with
2 introduction of a segment filter 4406.

3 More specifically, with reference to Fig. 45, render engine 222 identifies
4 multiple source clips from a common source which are not source time aligned,
5 block 4204 and/or require separate pre-processing filter(s), block 4206. Render
6 engine 222 generates a segment filter 4406 for the filter graph to reuse the source
7 and at least the source filter, block 4502. That is, the render engine 222 inserts a
8 segment filter 4406 between the source filter and one or more processing chains to
9 selectively provide otherwise disparate source clips from a single source. But for
10 use of the segment filter 4406 in the filter graph, the method 4500 of Fig. 45
11 executes in a fashion similar to Fig. 42, above.

12 Turning now to Fig. 46, a flow chart of an example method of segment
13 filter operation is presented, in accordance with one example implementation of
14 this aspect of the present invention. In accordance with the illustrated example
15 embodiment of Fig. 46, the method begins in block 4602, wherein the segment
16 filter 4406 seeks the source to the place that source data is first needed. As
17 introduced above, segment filter 4406 receives a request for source content from
18 matrix switch filter 308. It should be appreciated that insofar as segment filter
19 4406 may well support a plurality of processing chains coupled to a plurality of
20 matrix switch filters 308, a number of such requests may be received
21 simultaneously. According to one implementation, each of the matrix switch
22 filters 308 assigns a priority to the request for source content, wherein the priority
23 of the request changes as the time the content is needed draws near. According to
24 an alternate implementation, render engine 222 determines *a priori* whether source
25 content will be required simultaneously and, if so, provides a separate source chain

1 to accommodate such simultaneous content requests, thereby eliminating the
2 situation of the segment filter 4406 receiving simultaneous requests.

3 In block 4604, the source filter retrieves the requested content and passes
4 the data to the switch until some sort of indication is received that the end of
5 content has been received (e.g., an end-of-stream (EOS) indication, an application
6 interrupt, etc.). As introduced above, an application interrupt may be issued when
7 a user, through a user interface (e.g., media control application 224), wants to seek
8 to a certain point in the development project.

9 In block 4606, segment filter 4406 determines whether an EOS or an
10 application interrupt is received. If not, the process continues with block 4604. If
11 so, segment filter 4406 identifies the next required segment and when it will be
12 required, given the current seek location received from the matrix switch filter
13 308. Based, at least in part, on the current seek location, segment filter 4406
14 determines whether more segments of the source are required, block 4610. As
15 introduced above, if a user-defined seek command is issued, it may be issued to a
16 location in the development project where no further content is required from a
17 particular source. Thus, segment filter 4406 determines whether additional
18 segments are required in block 4610.

19 If no further segments are required along one of the processing chains
20 leading from the segment filter, render engine 222 may remove (at least
21 temporarily) that chain from the filter graph to free memory space and a matrix
22 switch filter input for other processing chains, block 4612.

23 If, in block 4610 further segments are required, segment filter 4406 issues a
24 seek instruction directing the source filter to retrieve and deliver the next segment,
25

1 in accordance with the matrix switch filter instructions, block 4614. This process
2 continues in an iterative fashion with block 4604.

3 Although the invention has been described in language specific to structural
4 features and/or methodological steps, it is to be understood that the invention
5 defined in the appended claims is not necessarily limited to the specific features or
6 steps described. Rather, the specific features and steps are disclosed as preferred
7 forms of implementing the claimed invention.
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25